



Final Report for ARO Research Agreement
Number DAAG55-98-1-0483, AMSSB-ACR
P-35631-MA-YIP

Sven J. Dickinson and Suzanne Stevenson

List of Manuscripts

A manuscript is in preparation, entitled, "Automatic Model Acquisition by Shape Abstraction" (Keselman and Dickinson), which will be submitted to ICCV and/or CVPR 2001.

Scientific Personnel

During the year covered by the awarded grant, one computer science graduate student, Yakov Keselman, was supported. The PI (Sven Dickinson) and co-PI (Suzanne Stevenson) are supervising his work under the grant.

Report of Inventions

There were no inventions to report during this period.

Scientific Progress and Accomplishments

Yakov Keselman, the student being supported by the grant, is addressing three important problems in computer vision.

1. In the domain of content-based image retrieval, the first problem deals with how to represent a query image in terms of a set of viewpoint-invariant indexing primitives, without knowledge of the object's identity or 3-D shape and without access to other views of the object. Our

approach, as laid out in the proposal and described in a paper presented in early 1998, relies on recovering the local 3-D part structure of the query object. Instead of exploiting knowledge of an object's structure (which is rarely known a priori), we will exploit knowledge of a finite, albeit large, vocabulary of 3-D parts that can be used to model a significant portion of any query object.

Yakov is nearing completion of a prototype system for modeling this universal part vocabulary, not only in terms of the 3-D shape of the parts, but their appearance in the image as well. His modeling system will allow us to build our part model representation used to identify user-outlined parts in the query image.

2. Once the parts are extracted from the query image, their possible views are sent as parallel (or ranked) queries to the image database. Before these can be matched to parts appearing in database images, the parts must be segmented from the database images. This represents perhaps the most challenging problem facing computer vision. How, from a plethora of structural detail, meaningless surface markings, and segmentation artifacts, can the coarse structure of a part be recovered from an image? Yakov is pursuing a novel approach to this problem, and it is this problem that Yakov has been focused primarily on over the last year.

We have developed an algorithm that will generate a class exemplar from a set of object exemplars presented to the camera in similar viewpoints. Each instance defines a lattice of possible region splits and merges. One possible path through this lattice defines an abstraction (or grouping) that is similar to the abstractions of the other instances. Yakov's algorithm searches for corresponding nodes (one per lattice) that represent the "most complex common denominator" across all the examples. We are in the process of testing the algorithm on both synthetic and real images, and a manuscript is being prepared describing the results (to be submitted to ICCV and/or CVPR later this Fall).

3. Once the abstraction is found, it defines a path in each exemplar's lattice which takes an image into the prototype. Moreover, each path defines a set of region split/merge training examples which will be input to a learning program. The goal will be to learn the conditions under which two adjacent regions should be split or merged, leading

to a set of perceptual grouping rules that can take an initial region map to a set of metaregions that correspond to the abstract surfaces on a part. These metaregions will be then matched to the part queries to perform effective indexing. Yakov has already performed some machine learning experiments to "tease out" the relevant parameters and region split/merge conditions.

Technology Transfer

There was no technology transfer to report during this period.